

# CCS News Bytes

LALP-05-015

November-December 2005



**People crowd in among the boxes as the Los Alamos National Laboratory booth at SC|05 takes shape.**

## CCS Had Major Presence at Supercomputing 2005

SC|05, the annual supercomputing conference held November 12-19 at the Washington State Convention and Trade Center in Seattle, Washington, offered an intense week of presentations, papers, and posters rich in new ideas and innovative developments.

It was an opportunity for leaders in high-performance computing to demonstrate their latest work and for colleagues and students to take a giant leap forward in knowledge—all in the course of just eight days.

The conference presented everything from razzle-dazzle to the most technical of scientific papers.

Microsoft Chairman Bill Gates used his keynote speech to launch a new product—Windows Compute Cluster Server 2003. Microsoft sponsored a “launch party” that was open to everyone in attendance and provided free food and drinks, a rollercoaster, games, and entertainment by singer Sheryl Crow.

IBM paid for a lobster dinner.

SGI and CRAY also sponsored parties.

But the “fun” was really a minor part of SC|05. There were dozens of tutorials, technical-paper presentations, poster presentations, panels, and

workshops—just to name a few of the activities available to participants.

In addition to Gates, the list of invited speakers in the conference program included Thomas J. Lange, director of modeling and simulation at Procter and Gamble; Laura Kong, director of the International Tsunami Information Centre; and Thomas Paul D’Agostino, acting deputy administrator for Defense Programs in the Department of Energy’s National Nuclear Security Administration. D’Agostino—the man assigned to make the final decision on who will manage Los Alamos National Laboratory (LANL) starting in 2006—spoke on “National Nuclear Security Administration Advanced Simulation and Computing: Ushering in a New Decade of Predictive Capability.”

More than 9,000 people attended the conference, and technical staff members from LANL’s Computer and Computational Sciences Division (CCS) were in the thick of it.

Becky Fernandez, executive office administrator in the CCS Division Office, compiled a list of 55 LANL people entitled to conference T-shirts—and 30 of them were from CCS.

**(Please see SC|05, page 2.)**

## List of R&D 100 entries includes three from CCS

The Computer and Computational Sciences Division has a proud history in RD100 competition. In 2003, work involving CCS claimed three RD100 awards.

The division’s winning entries in 2003 (all from CCS-1) were Clustermatic, the 10-Gigabit Ethernet Adapter (with Intel), and mpiBlast.

This year, the division will try to repeat its achievement.

Three of the 18 RD100 contest entries now in preparation at Los Alamos National Laboratory involve CCS staff members.

**(Please see RD100, page 2.)**

## RD100 (Cont'd from p.1)

- **RELOCATE (Reconfigured Logic Accelerated Traffic Engineering, formerly titled TRANSIMS-XLP)** was submitted by Justin Tripp, Anders Hansson, and Maya Gokhale. Tripp and Gokhale are in CCS-1. Hansson is in CCS-5.

TRANSIMS is a traffic simulator developed in CCS-5. In the past, it has provided simulations of the traffic in such cities as Portland and Los Angeles—but simulating one Portland day took about 16 hours on a 60-node cluster, and simulating one day in LA took about seven days on the same cluster. The larger the city, the longer the time, and fast action is vital for CCS-5 and its customers, who may be dealing with disasters.

RELOCATE uses field programmable gate arrays (FPGAs) to speed up the process.

FPGAs are integrated circuits—combining logic and memory—that can process digital information. They can be reprogrammed after their manufacture, and they can map algorithms directly into programmable logic.

The new FPGA system makes it possible to reduce the number of nodes required by TRANSIMS simulations from 60 to 12 while achieving the same performance. It also makes it possible to achieve speeds in TRANSIMS that are five times faster than in the past. The result: Traffic in Portland can be simulated in 3.2 hours, and traffic in LA can be simulated in 1.25 days.

- **The Trident Compiler for Supercomputing Application Acceleration** was submitted by Tripp, Kris Peterson, Jeff Poznanovic, Christine Ahrens, and Gokhale. Tripp, Peterson, and Gokhale are in CCS-1. Poznanovic was a student employee in CCS-1 in 2004 and part of 2005. Ahrens is in CCN-12.

In their presentation to the Laboratory entry-evaluation team, Trident team members described Trident as a C/Fortran compiler for floating-point algorithms that translates algorithms into reconfigurable logic.

Trident exploits the parallelism available in the input description. It also provides an open framework for experimentation, analysis, and optimization of floating-point algorithms on FPGAs and the flexibility to

integrate custom floating-point libraries easily.

Because Trident uses standard C input, it allows a larger number of application developers to add reconfigurable accelerators to their software.

- **EnergyFit™** was submitted by Chung-Hsing Hsu and Wu-Chun Feng of CCS-1.

Asked to describe their entry, Feng said, “EnergyFit™ is a system software package that intelligently tunes the frequency and voltage of processors (CPUs) in order to reduce energy consumption significantly while maintaining performance.”

“Specifically,” he added, “the software runs transparently on virtually any platform and delivers as much as 60% CPU energy reduction (20% on average) with little effect on performance. These energy savings translate into significant cost savings in datacenter and supercomputing environments, as for every watt of power that is saved in powering a system, another watt of power is saved from having to cool the system.”

He said, “Other EnergyFit™ benefits include improved system reliability (due to reduced component temperature) and reduced capital costs (for room engineering, climate-control systems, and replacement hardware).”

○○○

## SC|05 (Cont'd from p.1)

Gina Fisk from CCS-1 went to Seattle a week early because she was in charge of Booth No. 312, the booth for LANL. She had to oversee the provision of power and network, the building of the booth, and the installation of computers. “It was up for three days,” she said, “and then we had to tear it all down.”

It was “a big booth,” she said, “30 by 40 feet, with a 16-foot tower displaying the Laboratory logo.” Inside were a number of pods, each with a 12-foot-high tower of its own.

Among the topics represented in the pods were CCS-3’s Performance Architecture Laboratory (PAL); CCS-1’s RADIANT; Rod Linn’s EES-2 fire-modeling work; cluster advances and advanced architecture, featuring work by Maya Gokhale, Ron Minnich, and Greg Watson of CCS-1; CCN-5 network security research; and Carolyn Connor Davenport’s CCN-9 Network Processor Unit work.

## More SC|05 (Cont'd from p.2)

Conference participants circulated through the booth, soaking up the information presented. A few dropped by to look at LANL's recruiting poster and write down its website address. One of the educational programs featured on the poster was the "Supercomputing Challenge."

"We had a very good turnout," Fisk said. "We had a pretty good location this year, which was nice."

LANL's booth was in good—and familiar—company. The list of exhibitors, which covered six pages in the conference program, was divided into industry and research categories. The Laboratory's "research" exhibit was in the portion of the list that also included exhibits on Oak Ridge National Laboratory (ORNL); Advanced Simulation and Computing (ASC), presenting work by Lawrence Livermore National Laboratory and Sandia National Laboratories as well as LANL; Pacific Northwest National Laboratory (PNNL); Argonne National Laboratory; Brookhaven National Laboratory; Idaho National Laboratory; Lawrence Berkeley National Laboratory; Fermilab; the National Aeronautics and Space Administration; and the National Center for Atmospheric Research (NCAR)—among many others.

There was even a booth for universities on the Rio Grande—the University of New Mexico, New Mexico Institute of Mining and Technology, New Mexico State University, and the University of Texas-El Paso.

Fisk said the conference was a lot of work from her point of view but was worthwhile. Her favorite feature in the "fun" category was Family Day, November 16, on which participants were allowed to bring their spouses and children to the Trade Center and show them the conference exhibition. She gave her 21-month-old twins a tour.

An examination of listings in the conference program produced a substantial list of LANL participants.

### **Tutorials, Technical Papers, and Posters**

The SC|05 tutorial program was designed to offer "a compelling lineup of popular subject areas from previous years and new topics aimed at keeping participants at the forefront of HPC (high-performance computing)." The program said that, "As in the past, the half-day and full-day sessions were chosen from a competitive field of submitted proposals."

The list of presenters for one of the tutorials, "M02: Component Software for High-Performance Computing; Using the Common Component Architecture," included Craig E. Rasmussen of CCS-1. The two men listed as presenting the tutorial entitled "M03: A Practical Approach to Performance Analysis and Modeling of Large-Scale Systems" were Darren J. Kerbyson and Adolfo Hoisie, both from CCS-3.

The "Technical Papers" portion of the conference program said SC|05 would "offer 62 outstanding papers covering a wide range of research areas and presenting a truly international perspective on high-performance computing and networking. These papers report experimental and theoretical results, innovative hardware and software design, and related case studies that advance the state of the art and bring together the latest achievements in technology, research, science applications and education."

Chung-Hsing Hsu and Wu-Chun Feng, both of CCS-1, presented one of the papers, entitled "A Power-Aware Run-Time System for High-Performance Computing." Roberto Gioiosa of Rome University in Italy, Jose Carlos Sancho, Song Jiang, and Kei Davis of CCS-3, and Fabrizio Petrini of PNNL, formerly of CCS-3, presented another, entitled "Transparent Incremental Checkpointing at Kernel Level: A Foundation for Fault Tolerance for Parallel Computers." A paper entitled "On the Feasibility of Optical Circuit Switching for High-Performance Computing Systems," included three Los Alamos names on its list of authors: Kevin J. Barker, Hoisie, and Kerbyson, all of CCS-3. And three of the four names listed on the paper "Partitioning Hardware and Software for Reconfigurable Supercomputing Applications: A Case Study," were from LANL: Justin L. Tripp and Maya Gokhale are in CCS-1, and Anders A. Hansson is in CCS-5.

LANL people chaired three of the technical-paper sessions. Andy White of the CCS Division Office was the chairman for the session on "High End Systems Issues"; Scott Pakin, CCS-3, chaired the session on "MPI and Network Transport"; and Feng, CCS-1, chaired the session on "Power and Programmability."

The poster entitled "PFLOTRAN: A Massively Parallel Simulator for Reactive Flows in Geologic Media" listed three names, including two from Los Alamos: Chuan Lu and Peter C. Lichtner, both of EES-6.

**(Please see Remaining SC|05, page 4.)**

## Remaining SC|05 (Cont'd from p.3)



One of the LANL booth “pods” in use.

### Panels, Award Sessions, ‘BoF’ Meetings, and Several Special Initiatives

Feng moderated two panels—one on “The Six-Million Processor System,” and one entitled “Tour de HPCycles.”

Hoisie was the chairman for two awards sessions.

Timothy C. Germann of X-7, Kai Kadau of T-14, and Peter S. Lomdahl, also of T-14, presented “25 Tflop/s Multibillion-Atom Molecular Dynamics Simulations and Visualization/Analysis on BlueGene/L” during the Gordon Bell Prize Presentations.

During one of the “Birds-of-a-Feather (BoF) sessions,” called to discuss topics of mutual interest to a group of participants, Watson and Rasmussen, both of CCS-1, gave a presentation entitled “Using Eclipse to Develop Parallel Applications.”

Cindy Sievers of CCS-1 was among those whose names were listed on a presentation entitled “Models for Collaborative Projects in Emerging Advance Computing Communities.” It was part of an “SC Global” session on “Developing Remote Collaboration.”

One of the finalists in Session 1 of the “HPC Analytics Challenge” was entitled “Network Traffic Analysis with Query-Driven Visualization.” It bore an extensive list of names that included Mike Fisk, CCN-5, Paul Weber and Steve Smith, D-4, and Chris Davis, NMT-DO.

The name of Parks Fields, CCN-9, was one of two on a StorCloud presentation entitled “Visualization of High Resolution Simulation of a Compressible, Turbulent Flow.”

The names of Feng, Mark Gardner, and Jeremy Archuleta, all CCS-1, Heshan Lin (listed as

LANL/North Carolina State University), and Venkatram Vishwanath (LANL/University of Illinois-Chicago) were among those listed for “mpiBLAST on the GreenGene Distributed Supercomputer: Sequencing the NT Database against the NT Database (an NT-Complete Problem),” an entry in the StorCloud Challenge.

### Special Thanks

Twelve LANL people were among those thanked in the last section of the conference program for their contributions. Here is the list (with complete names): Adolfo Hoisie, CCS-3, Awards chairman; Harvey Wasserman, CCN-7, Scatter/Gather chairman; Corinne Fresquez, CCN-18, Conference Office; Cindy Sievers, CCS-1, general producer of SC Global Showcase; Park Fields, CCN-9, for his work on “Open Infiniband”; Denny Rice, CCN-5, for his work on Power Engineering and for his work as lead for Vendor Support; Wu Feng, CCS-1, Technical Paper Committee chairman for the Networking Area; Mark K. Gardner, CCS-1, Chung-Hsing Hsu, CCS-1, Scott Pakin, CCS-3, and Andy White, CCS-DO, all members of the Technical Paper Committee; and Henry John Alme, X-8, a member of the Poster Committee.

(Please see SC|05 Photos, page 5.)

ΩΩΩ

### *Urbatsch Goes Recruiting*

Todd Urbatsch, group leader of CCS-4, has been doing substantial recruiting recently. He visited both The Pennsylvania State University and the University of New Mexico in search of promising students and postdoctoral employees.

Asked about the response he received, he said it was good. “We’ve got leads on four or five summer students or postdocs or future employees,” he said.

His efforts have been “fairly informal.” He did his recruiting in conjunction with accepting invitations from nuclear engineering departments to come and speak. He took along one-page recruiting posters to hand out or tack to bulletin boards, mini-posters of the CCS-4 panel that resides outside the main auditorium, and copies of CCS News Bytes. He found, in some cases, that students he met were already familiar with CCS-4 because they had perused its website.

He warned, however, that sometimes, previous connections can cause complications. Urbatsch went to graduate school at the University of Michigan—and he spoke in Pennsylvania right after Michigan defeated Penn State in football.

**SC|05**

At left, empty space at the Washington State Convention and Trade Center in Seattle awaits work on the Laboratory's booth. Below, the LANL booth takes shape.



Bottom left, the familiar Laboratory logos appear. Below right, the completed booth attracts users.



Special thanks to Gina Fisk for use of her photos on this page and on pages 1 and 4.

## Meet Our New People



**Fidela Gonzales**

There's a new face in the Computer and Computational Sciences Division Office (CCS-DO).

**Fidela Gonzales** joined the division on October 24 as an administrative specialist.

Before joining CCS-DO, Gonzales had been employed by the University of California for about 18 months.

She worked in Stockpile Complex Modeling and Analysis (D-2), Quality Systems Support and Production Control (MSM-2), the Nuclear Nonproliferation Division Office (N-DO), and (very briefly) in Nuclear Materials Technology (NMT-11).

Her background includes three years as a contractor with the Plus Group and substantial experience with the State of New Mexico. She was an eligibility worker for Human Services in Española and Taos, and she worked as a records specialist with the Statistics Bureau in Santa Fe.

She is currently attending Northern New Mexico Community College in Española, seeking a degree in business administration.

Gonzales and her husband, Orlando, have two sons—James, 26, and Jason, 16. James and his wife, Jenelle, have a 20-month-old daughter, Jenique, and a second child on the way. Jason is an 11<sup>th</sup> grader at Peñasco High School.

Asked why she moved from N-DO to CCS-DO, Gonzales said, "I saw this job advertised. It was a promotion for me."

ΩΩΩ

**Timothy Kelley** joined the Transport Methods Group (CCS-4) September 26 as a technical staff member. He is working on the JAYENNE Project.

Kelley was born in Chicago, but New Mexico is not a new experience for him. He earned a

bachelor's degree in liberal arts at St. John's College in Santa Fe in 1994.

He also spent several years doing post-graduate research at the Los Alamos Neutron Science Center (LANSCE).

He earned a master-of-arts degree in physics from Brandeis University in Waltham, Massachusetts, in 1996, and he also holds a doctorate in physics, which he earned at the University of California-Riverside in 2002.

After he completed his doctorate, he worked at California Institute of Technology in Pasadena, doing software engineering, until he responded to a job advertisement and came to Los Alamos National Laboratory (LANL).

Asked what attracted him to LANL, he cited "the scope of research here." He said, "There's something here for everyone ... I'm interested in developing excellent software for scientists to use. There's plenty of scope to do that here."

Another reason for his interest in the Laboratory was more personal. His wife, Karen Kelley, who has a doctorate in nuclear engineering from Georgia Institute of Technology, works in Engineering Sciences and Applications-Weapon Response (ESA-WR).

"We met at LANSCE," Kelley said. They dated in Los Alamos, maintained a long-distance relationship when he left, and married two years ago. Their first child, Lucy, was born in Los Alamos on October 27.

ΩΩΩ



**Timothy Kelley**

## Security Tip

The hustle and bustle of the holiday season makes it more likely that Laboratory workers will be distracted, but employees at Los Alamos National Laboratory (LANL) work in an environment where security is vital.

We all need to focus and follow the rules.

Don't forget that LANL security lists as "prohibited articles" (articles "not permitted at the Laboratory unless approved in advance by S Division's Office of Security Inquiries") the following items: "any dangerous weapons," "alcoholic beverages" and "controlled substances."

Always take a second to think before you enter a security area (an area behind the fence at TA-3, for example). Read the signs and be sure what you can and cannot bring into the area.

Remember to wear your badge in a visible position above your waist, and always check to make sure that it is securely attached after you have used it to go through a gate. Don't lose it!

## In Case You Missed It....

Two well-known employees have departed from CCS Division recently, accepting other positions.

**John Ziebarth**, group leader in CCS-1, left to become senior vice president in the Krell Institute in Ames, Iowa, which lists as its goal an "active partnership with the science, technology, and education communities."

**David Carrington** has left CCS-4, but his new office will be closer to old friends. He has joined the Computational Fluid Dynamics Group (T-3) and will be located in Building 200.

**Kevin Roark in Public Affairs recently sent out this news release (edited lightly for style and space considerations in this publication):**

LOS ALAMOS, N.M.—Award-winning Los Alamos National Laboratory-developed software is helping researchers here and elsewhere better understand a database of biological information and enable a plethora of biological studies from organism "barcoding" to gene function and evolution.

The software, mpiBLAST, coupled with a supercomputer assembled over a high-speed network and distributed across the country just for this purpose, will make the biological information stored in large databases more useful for researchers by enabling a Google-like indexing structure that tracks relationships among the sequences in these large databases. Such an indexing structure can increase search speed times by a factor of 100 while at the same time providing an up to 20-fold compression in the size of the database.

mpiBLAST, an open-source project led by Los Alamos researcher Wu Feng, is being tapped to

harvest the NT biological database in order to create the Google-like indexing structure. Los Alamos researchers announced at the Supecomputing 2005 Conference that they will lead a large-scale nationwide effort to sequence-search the entire NT database.

The NT biological database is akin to a "biological dictionary organized as a flat file." When biologists need to know if a particular genomic sequence has already been catalogued, they look through this dictionary for that genomic sequence. If they can't find the desired sequence, they add the new information to the end of the file thus making the unordered file larger and larger.

With the idea that it would be much better to organize the database and build it with some structure that is searchable in a nonlinear manner, Feng and other scientists, using the "GreenGene" supercomputer, intend to give this huge database that structure by sequencing the entire database against itself.

"If this endeavor to sequence-search the entire NT database succeeds, the result of this experiment will provide critical information to the biology community, including insightful evolutionary, structural, and functional relationships between every sequence and family in the NT database," notes Feng, principal investigator for the project. In all, the large-scale experiment is expected to generate 100 terabytes of output—enough to fill-up roughly 2,000 iPods.

mpiBLAST, as distributed by the Laboratory (<http://mpiblast.lanl.gov> or <http://www.mpiblast.org>), won an R&D 100 Award in 2004. It is a search tool that enables biologists to characterize an unknown sequence by comparing it against a database of known sequences. The similarity between sequences then enables biologists to detect evolutionary relationships and infer biological properties of the unknown sequence. On a 128-processor supercomputing cluster, mpiBLAST can deliver a speed-up of 305-fold, thus decreasing the search time of a representative 300-kilobyte query file from nearly 24 hours down to only 5 minutes. Additional speed-up, as provided by a parallelized I/O version of mpiBLAST called mpiBLAST-pio, reduces the search time further and allows the code to scale to larger system configurations.

Led by the Laboratory, the nationwide team working together on this collaborative endeavor includes industrial participants from Intel, Panta Systems, and Foundry Networks. Academic and government participants are from North Carolina State University, Oak Ridge National Laboratory, Utah, and Virginia Tech universities.

ΩΩΩ

**CCS News Bytes wishes its readers a happy holiday season and a productive and satisfying New Year.**

## Closing Out 2005

CCS Division Leader William Feiereisen (right) closed out 2005 at an all-hands meeting, summarizing the major accomplishments in all five CCS groups. Afterward, employees enjoyed an elegantly presented buffet (below, left and right). Associate Director for Weapons Physics Susan Seestrom came by to join the party, providing the latest news on the Laboratory transition (right, third photo down). She presented coffee cups to all CCS employees (bottom right), and she posed with Feiereisen (bottom left) for a holiday picture.



Photos by Charmian Schaller